Performance Studies of Out-of-Core Computing on Various Storage Devices – Fusion-IO's IOSan, SuperTalent's RaidDrive, OCZ Vertex 3, Kingston V+

Derrick Lee (UNCC), Garrett Lord (MTU), Jesse Martinez (NMT), Christopher Moore (UK)
Mentors: HB Chen, Parks Fields, Benjamin McClelland, Alfred Torrez
Computer Systems, Cluster, and Networking Summer Institute 2011

- Datasets are characterized by their very large sizes ranging from hundreds of gigabytes to hundreds of terabytes with multiple superposed scalar and vector fields, demanding an imperative need for new interactive exploratory visualization capabilities
- Algorithms must be optimized to efficiently fetch and access data stored in slow bulk memory such as hard drives, solid-state-devices (SSD), or tape drives
- Sequence of Out-of-core benchmarking tests were done on various storage devices such as SATA based Solid State Devices and PCI-Express based SSDs
- Devices Tested: Fusion IO 640 GB, Super Talent RAID Drive 512 GB, 4x 240 GB OCZ Vertex 3, 4x Kingston V+ Series 256 GB

FIO Benchmark Baseline

The following graphs are the FIO output results for random read/write and sequential read/write tests reporting the I/O operations per second (IOPS) metric and aggregate bandwidth. Each of the graphs shown report IOPS and aggregate bandwidth for their respective SSD for un-buffered results at 64 jobs.

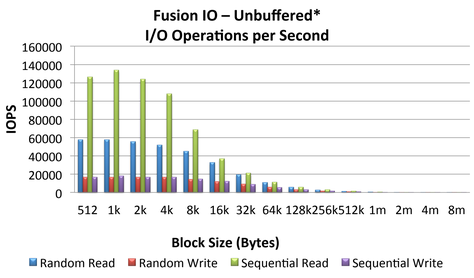

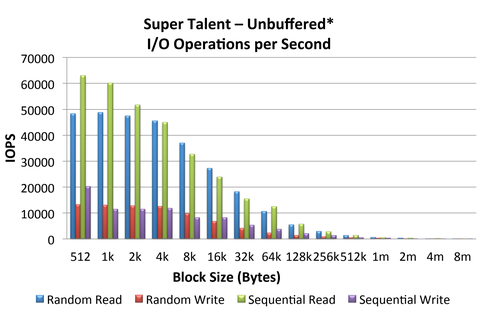

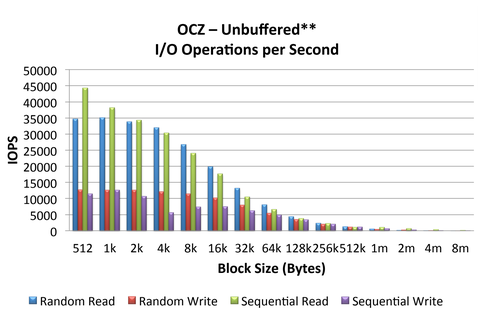

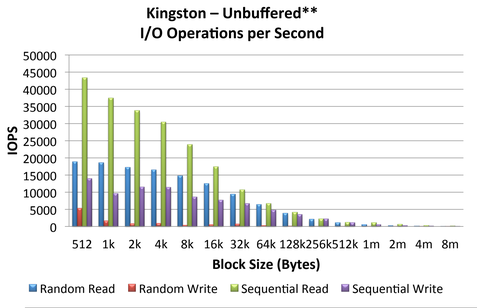

- With a lower block size and larger number of jobs being performed, the reported IOPS (left) appears to be at max range.
- For Aggregate Bandwidth (right), there is a direct correlation with a higher block size and number of jobs for each device.
- 2k-8k block sizes showed to be the optimal range for each of these devices.

Device*** (MTBF)	Capacity (GB) / Dollar	Bandwidth (MB) / Dollar	IOPS/ Dollar
Fusion IO	0.040	0.012	3.011
Super Talent (1.5 mil. Hours)	0.256	0.056	14.355
OCZ Vertex 3 (2 mil. Hours each)	0.426	0.035	8.904
Kingston V+ (1 mil. Hours each)	0.499	0.021	7.224

*Fusion IO and Super Talent are using ext2 file system
**OCZ and Kingston are using ext2 file system with RAID 0 at 128k stripe size and read/write cache
***Bandwidth and IOPS are calculated with the AVERAGE results of all tests with 4k block size at 64 jobs divided by the costs

Note: Graph units are not to scale

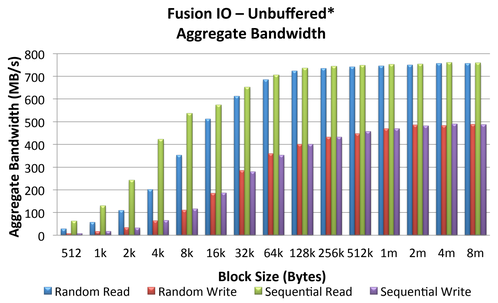

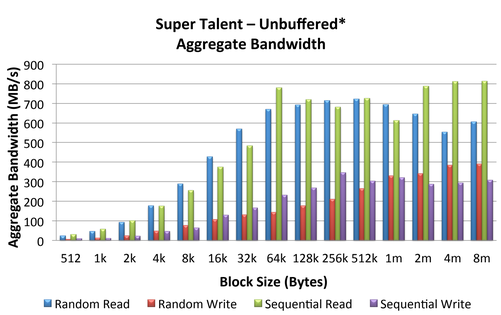

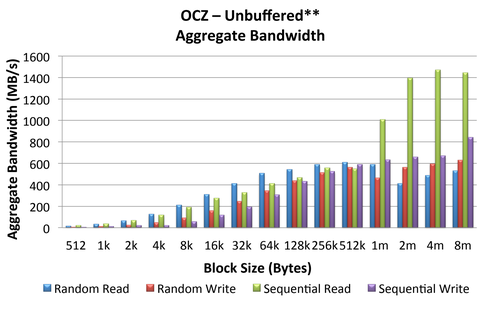

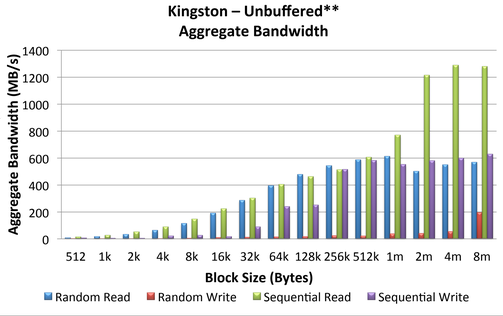

Graph 500

The Graph 500 benchmark was used to analyze memory SSD performance by generating large scale graphs to traverse the edges between the SSD and the memory. The following graph displays the TEPS (Traversed Edges per Second) for the Fusion IO SSD and memory baseline test. Additionally, the minimum, median, and maximum TEPS are reported for each case.

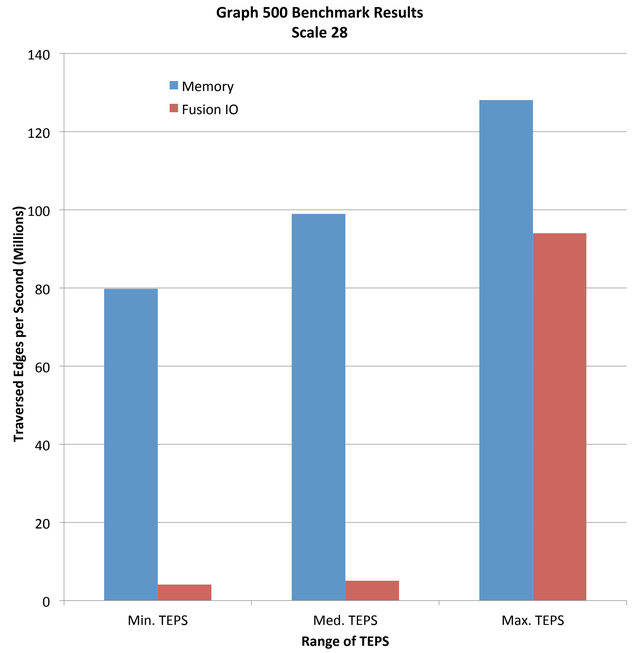

- Memory testing involved 144GB of RAM in which 128GB was used by Scale 28 Graph 500 standards.
- Fusion IO testing had 94GB of RAM with 70GB ext2 swap on the device.
- Future Tests will be done on the other SSDs.

TEPS (Millions)	Memory	Fusion IO
Minimum	80	4
Median	99	5
Maximum	128	94

Conclusion

- In respect to processing power, FIO test results also showed that the Fusion IO drive displayed significantly better performance. However, performance per dollar indicates that the smaller SSDs had a higher cost efficiency rating.
- Graph500 benchmark testing showed that the Fusion IO had a high edge traverse rate with the swap system but baseline results showed to be higher.
- Although the Fusion IO had higher performances, the performance/cost results showed to be lower than the rest of the SSD devices.
- Contact HB Chen at (hbchen@lanl.gov) for future inquiries

LA-UR-11-11208